NComputing L-Series LAN Deployment

Best Practices for Local Area Network Infrastructure

Scope:

NComputing's L-Series terminals connect to a host computer through an Ethernet interface and IP protocol. The L-Series terminal relies on a robust network infrastructure, with minimal latencies for correct operation and a satisfactory user experience. This document explains the minimum network requirements and makes recommendations for a successful deployment of a large number of L-Series terminals on a single site.

Audience:

The target audience for this document includes IT personnel (ultimately those individuals in charge of the network infrastructure where the deployment is to happen), Technical Sales, resellers, VAR's, and other people with some background in IP networks.

Background:

To better understand the conclusions of this white paper, let's recap the functionality of a few basic elements of a Local Area Network (LAN):

> **Router(s):** Equipment that interconnects and route IP packets between distinct networks, defined by IP address ranges. It is usually the connection point between the Local Area Network (LAN) and the outside world (the Internet or another Wide Area Network, WAN). Routers use the destination IP address of the network packet in order to choose the proper forwarding interface for it. Routing rules (or Routing Tables) are generally fixed (pre-programmed for a specific network topology) but they can adjust their internal tables based on information passed along by other routers located upstream or downstream from them. Routers can also provide other network functions as well such as NAT (network address translation), Firewall, and DHCP server.
>
> **Switch(es):** Equipment where all network terminal devices (computers, L-Series terminals, network printers, etc) connect to. Its main function is to aggregate all network traffic based on physical location, and forward the network packets according to the physical address of the destination device (MAC address). Unlike routers, switches are constantly learning what the best (fastest) path between network devices A and B is. They do so by collecting MAC address information for the devices that are connected to its ports and building tables with that information. For example: If a switch notices that on its port number 5 it only receives packets from a single device (say MAC address XYZ), it adds a table entry that basically says "all packets for device XYZ goes to port 5, and only port 5". After the network and its

connected devices are running for a while, the network switches should have figured out where all devices are located and therefore network packets start to flow very efficiently from source to destination device, without being un-necessarily sent to un-related devices on the network. Also, switches can implement Virtual LAN functionality (VLAN), which we'll discuss next.

VLAN(s): Sometimes it is advantageous to break a large network into smaller networks, grouping together devices that belong to the same "logical" network. For example, VLAN 22 groups all computers and devices that belong to the Accounting department and they will only communicate among themselves (lets say, for security reasons). VLANs are usually defined at the Switch level, on a port-by-port basis, and are static by nature. Any changes have to be done by the network administrator. The idea behind VLANs is that network packets are not allowed to propagate outside the switch ports defined by a VLAN, specially Broadcast packets (packets intended for ALL network devices), that, if allowed to run free on a large network, could easily saturate it (Ex: a single Windows machine sends out many broadcast packets per hour). A LAN could have up to four thousand VLANs, depending on network equipment vendor and model.

Inter VLAN communication: When network equipment on one VLAN needs to send / receive IP packets to / from devices located on another VLAN, these packets need to converge to a Gateway or Router, where they "jump" to the other VLAN. The number of devices a packet that started on let's say ***VLAN n,*** needs to go through in order to reach a machine on ***VALAN m*** depends on the network topology but it's definitely more than three (switch-router-switch). The latency (or delay) for a packet to travel from its origin to its destination depends also on the performance of the devices in-between. Except for high end routers, used mostly on metropolitan WANs and costing tens of thousands of dollars, mid-range routers were not designed to route data at "wire speed" and a considerable delay is expected when packets have to go through them. As an example, a mid range router from a well known manufacturer can route packets at a maximum of 50 Mega bits per second, far slower than a low cost Gigabit switch that can sustain a transfer data rate at the same speed as the port connection (100Mbit/sec in, 100Mbit/sec on each port)

L-Series Terminal network utilization:

NComputing has developed a very efficient Terminal Services protocol that runs on top of the TCP/UDP/IP protocol. Any and all screen updates on the L-series terminals monitor are converted into TCP packets and sent to the L terminal instantly. Likewise, every keystroke a user types on an L-series terminal attached keyboard is encapsulated into a TCP packet and sent to the computer running NCT Terminal Server software. The same is true for audio and USB communications between L-series terminal and Host Computer.

An L-series terminal can work well with as little as 200 Kbps of network bandwidth on a pure data entry usage model, but it could also scale to several Mega bits per second of required network bandwidth in a high end usage environment (video playback or USB storage data transfers to/from host computer).

In a typical office environment, where users type away at documents, crunch numbers on a spreadsheet or check e-mail and browse the web, the network should have allocated bandwidth requirement at per 1 to 1.5 mega bits per second per terminal (No USB transfers, no streaming Audio, no Video…).

In a high end usage model (streaming video and/or heavy USB transfers), the network should have allocated bandwidth requirements at per ~15 mega bits per second per terminal. In order not to saturate the LAN backbone and network equipments in between an L-series terminal and the Server Host computer, an L-series terminal limits the amount of bandwidth it consumes to a max of ~15 Mega bits per second, on each direction. This has many advantages but also some deployment needs that the network administrator should be made aware of before a large scale deployment.

NComputing operates with a USB 1.1 emulation model and transfers at speeds slightly slower than most 1.1 devices. It that also emulates some USB 2.0 functions for compatibility, thus maximum USB bandwidth over the LAN is less than 4 Mbps.

Video playback applications are also very demanding on the network, on the host and on the L-terminals. With that in mind, NComputing's Terminal Services protocol implements patent pending compression algorithms that allows very good video playback (ex:360 x 240 pixels @ 15 frames per second) without dropping frames. Bandwidth here can be up to 15Mbps. It is not advised to run both streaming video and large USB transfers at the same time.

Switches and Network backbone provisioning

Given the requirements from the previous topic, and assuming a worst case scenario, the rule of thumb is that ***for every Seven L-series terminals, the network backbone must provide one dedicated 100Mbps channel into the host server.*** If this condition is not met, the user experience may deteriorate. That is: USB memory drives will take much longer to transfer files to/from server, or video and audio playback may drop frames proportionally to available network bandwidth, etc.

If the available network bandwidth falls much below what is required for basic communications between the L-station and the host computer it is connected to, and network packets start to get dropped / delivered out of order, even keyboard and mouse information could be lost, causing the terminals to start an erratic behavior or shutdown the user session altogether.

Therefore it is important that when deploying more than 7 terminals in an active environment using video that a dedicated Gigabit backbone be used. In deployments that have a fixed application use that does not include video or USB then many more units can be supported on a 100Mb backbone.

In addition all terminals must be deployed on a single subnet that connects directly to the host server. This includes VLAN's that must be programmed for the same subnet.

Summary of Best Practices for Network Implementation

(For heavy traffic environments)

1) When subnets are used, make sure all terminals and the host are connected to the same VLAN and subnet.
2) Use at least one network card with gigabit capabilities on the Host computer to connect to the L-series terminals, future-proofing your network and to make sure more than 7 L-series terminals will work at full speed.
3) Likewise, if more than 7 L-series terminals are used in an active video environment, avoid using a switch that is limited to 10/100 Mbps.
4) Use a separate NIC card to connect the Host computer to the rest of the network and the internet. Bridging both NIC's in software is OK.
5) Whenever possible, use switches with gigabit backbone to isolate unrelated network traffic from the L-series terminals.

Examples:

Good provisioning: Ideally, the network connection between the terminals and the host server should be free of unrelated traffic. That is why two network cards per server is always a good thing to have (also good for fail over scenario).

a) One server equipped with two Gigabit Network adapters (NICs), one 24 or 48 port switch with at least one Gigabit port (regular ports could be 100Mbps) and 24 L-series terminals. Doing the math, a Gigabit NIC or switch port has enough bandwidth for about 55 L-series terminals running full speed, so this configuration would perform just fine, with room to grow for a second host. (Maximum 30 terminals per host.)

b) One server with two gigabit NICS, where one of them is connected to a 5-port Gigabit switch whose remaining 4 ports are connected to four 8-port 10/100 switches. The remaining 28 ports on the 10/100 switches could accommodate 28 L-series terminals running at optimum speed.

c) One server with two Gigabit NICs, where one of them is connected to a Gigabit capable port on a high end switch. This switch is partitioned into VLANs and interconnects to other switches through one or more gigabit connections. 30 L-series terminals are connected to 10/100 ports on other switches AND these ports are programmed as part of the SAME VLAN as the server they will connect to.

©2007 NComputing, inc.

Low cost, well provisioned 29 user solution

Host PC w/ Giga bit Ethernet card

5 port Gigabit switch
D-link DGS-2250

8 port 10/100 switch
D-link DES-1108

User 1 **User 2** **User 3** **User 7**

8 port 10/100 switch
D-link DES-1108

User 8 **User 9** **User 10** **User 14**

8 port 10/100 switch
D-link DES-1108

User 15 **User 16** **User 17** **User 21**

8 port 10/100 switch
D-link DES-1108

User 22 **User 23** **User 24** **User 28**

©2007 NComputing, inc.

Bad provisioning:

a) Same as a) above but the servers only have 10/100 NICs. (Only 7 L-series terminals are possible in this case; unless there is very light traffic generated by the terminals.)

b) Same as c) above but server and L-series terminals are on different VLANs. This is the worst possible case because all network packets to/from the L-series terminals would have to be forwarded to a network device (usually a Router or Gateway) where these packets could transpose VLANs and go back to their destination. In all probability, there will be a bandwidth bottleneck somewhere along the best path (which could easily have 4 or 6 hoops) which, when compounded to the Router/Gateway processing bandwidth capabilities, would create an un-acceptable user experience, or even unpredicted behavior.

c) Any condition where the network cannot guarantee at least 15Mbps of dedicated bandwidth, each way, to and from an L-series terminal.

Summary:

Best practices should always be used when provisioning the network infrastructure for large L-series terminal deployment to maximize the user experience. Enjoy your NComputing experience.

Example of an excellent deployment:

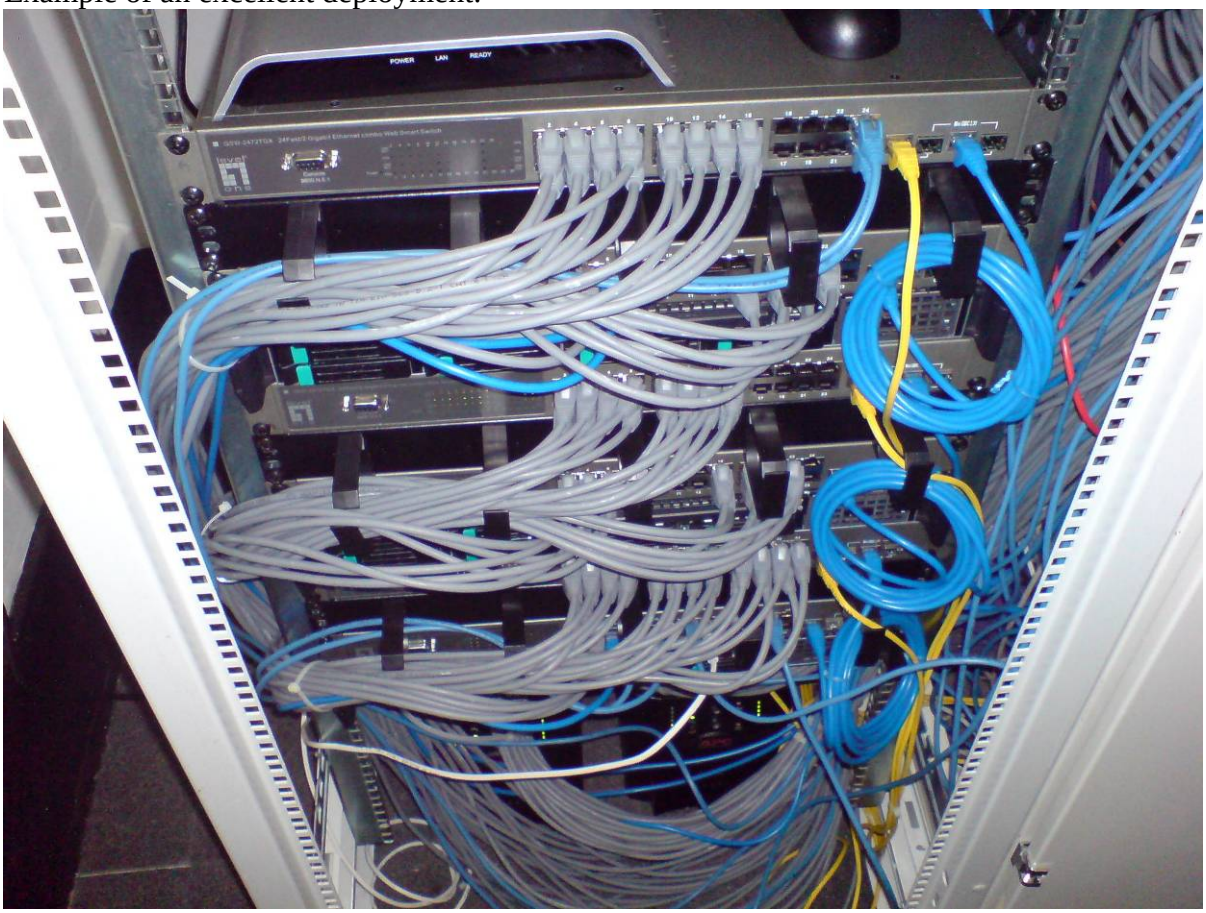

.

This multi 30 User environment was crated with a 1Gb backbone and uses color coded cables to indicate their function (gray cables to the L-Series terminals).